

V. Elements of Digital Electronics and Signal Processing

1. Elements of Digital Electronics

- Logic functions
- Logic circuitry
- Power dissipation
- Timing problems

2. Digitization of Pulse Height and Time – Analog to Digital Conversion

- Resolution
- Differential non-linearity
- Integral non-linearity
- Conversion time
- Count-rate performance
- Stability

Analog-to-digital conversion techniques

Time digitizers

3. Digital Signal Processing

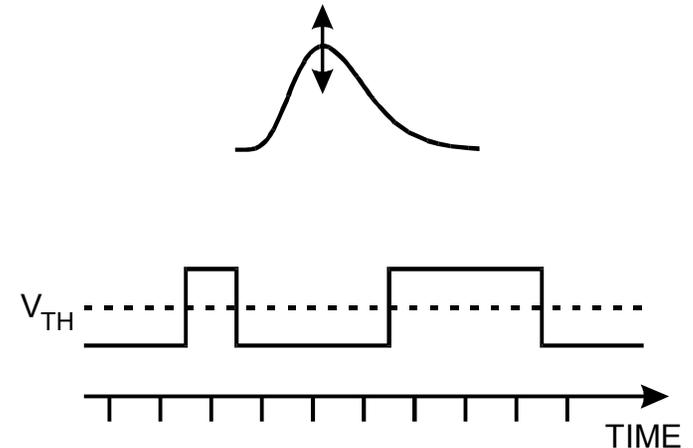
1. Elements of Digital Electronics

Basic differences

Analog signals have variable amplitude

Digital have constant amplitude, but variable timing

Presence of signal at specific times is evaluated:
(does the signal level exceed threshold?)



Transmission capacity of a digital link (bits per second)

Shannon's theorem:

$$C = B \log_2 \left(1 + \frac{S}{N} \right)$$

B = Bandwidth

S = Signal (pulse amplitude)

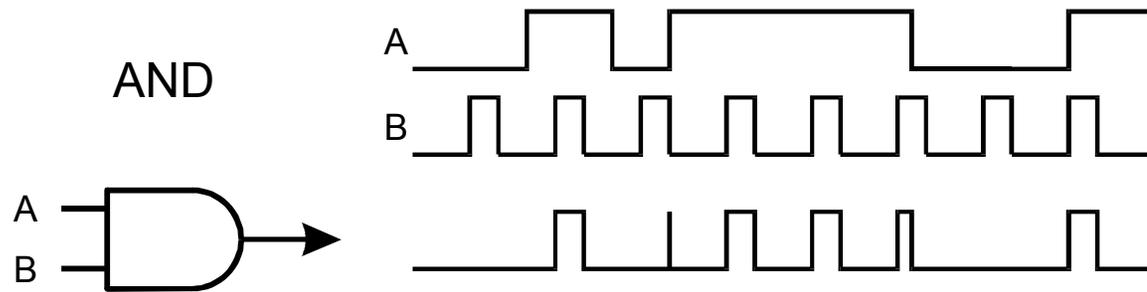
N = Noise

Noise enters, because near the switching threshold, digital elements are amplifiers.

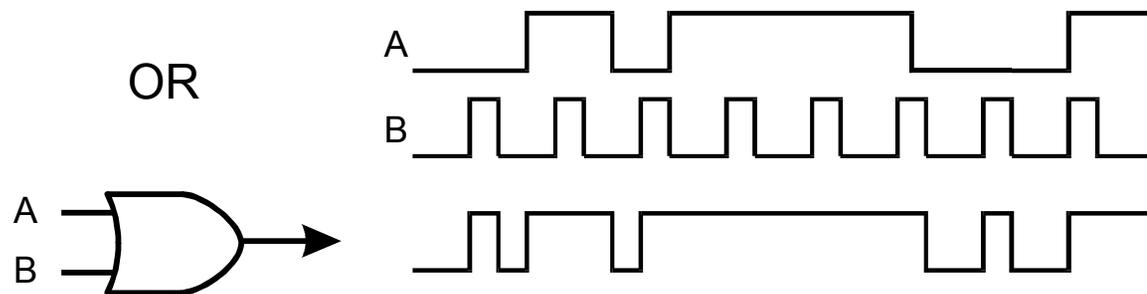
If the noise is due to cross-talk from other digital signals,
increasing the pulse amplitude will not improve S/N

Digital electronics not just a matter of “yes” or “no” – real systems must also deal with “maybe”.

Logic Functions



OUTPUT IS HIGH WHEN **BOTH** INPUTS ARE HIGH.

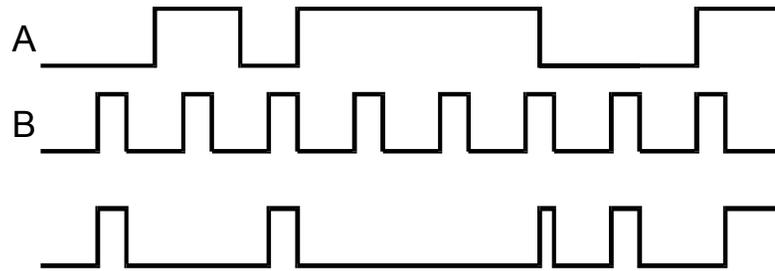
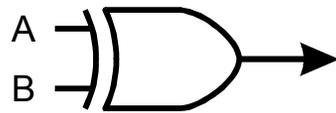


OUTPUT IS HIGH WHEN **ANY** INPUT IS HIGH.

When outputs inverted: NAND and NOR

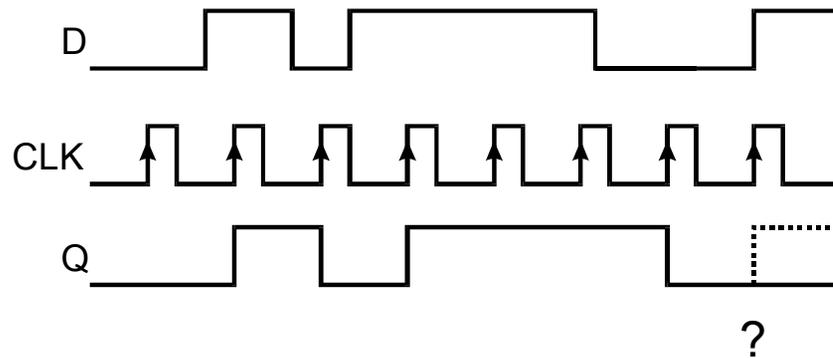
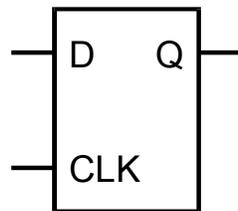
Logic Functions cont'd

EXCLUSIVE
OR



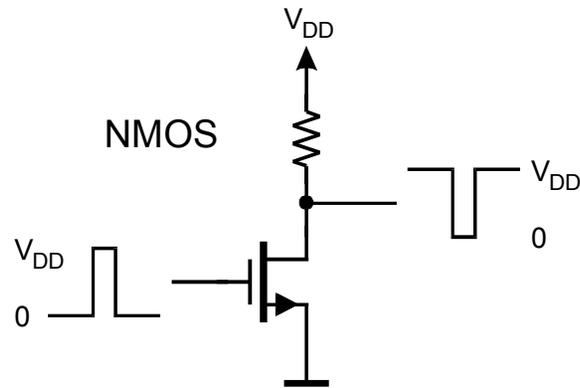
OUTPUT IS HIGH WHEN ANY **ONE** INPUT IS HIGH.

D FLIP FLOP
(LATCH)

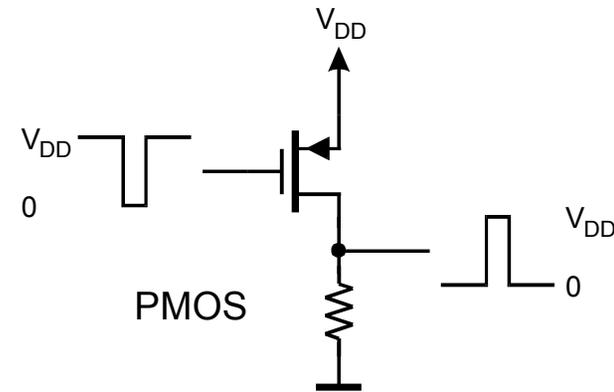


LEVEL AT D INPUT IS STORED AT POSITIVE CLOCK (CLK) TRANSITION.
 AMBIGUITY WHEN D AND CLK TRANSITIONS ARE SIMULTANEOUS.

LOGIC CIRCUITRY - Modern logic uses MOS technology.



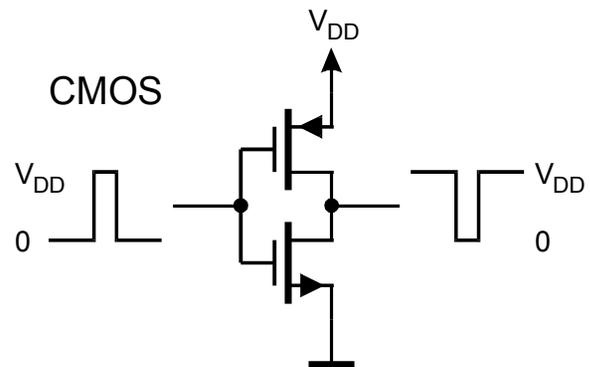
Transistor conducts when input is high.



PMOS: Transistor conducts when input is low.

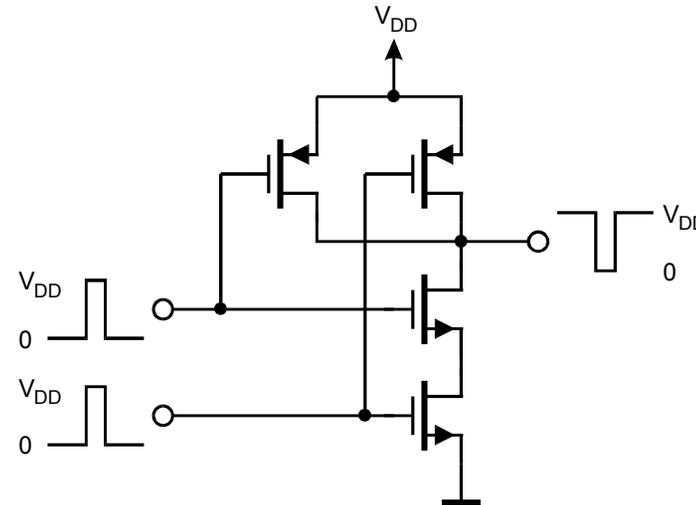
CMOS – combine NMOS and PMOS \Rightarrow significant power reduction

CMOS Inverter



Current flows only during transition.

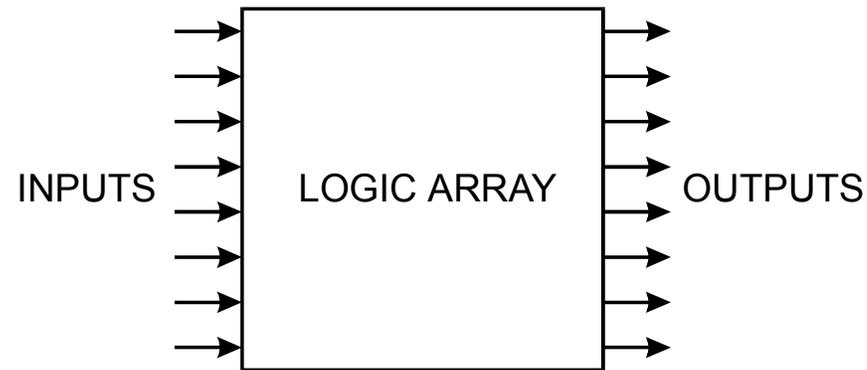
CMOS NAND Gate



LOGIC ARRAYS

Complex logic systems are not designed using individual gates.

Instead, logic functions are described in a high-level language (e.g. VHDL) and synthesized using design libraries (in custom ICs, “ASICs”) or programmable logic arrays.



Typical: 512 pads usable for inputs and outputs, $\sim 10^6$ gates, $\sim 100\text{K}$ memory

Software also generates “test vectors” that can be used to test finished parts.

POWER DISSIPATION AND PROPAGATION DELAYS

Energy dissipated in wiring resistance R :

$$E = \int i^2(t) R dt$$

$$i(t) = \frac{V}{R} \exp\left(-\frac{t}{RC}\right)$$

$$E = \frac{V^2}{R} \int_0^{\infty} \exp(-2t/RC) dt = \frac{1}{2} CV^2$$

If pulses (rising + falling edge transitions) occur at frequency f ,

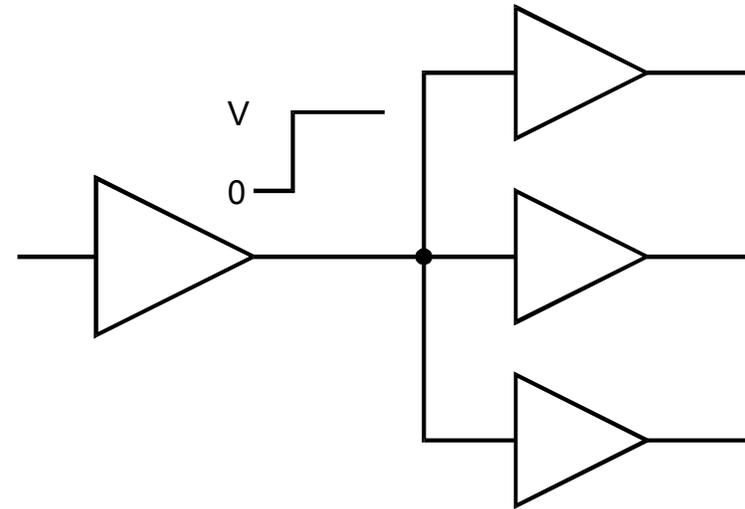
$$P = fCV^2$$

Power dissipation increases with clock frequency and (logic swing)².

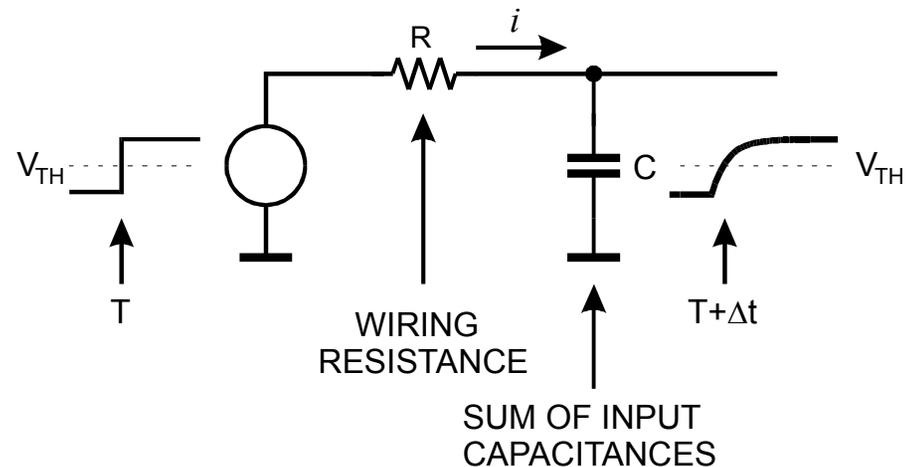
RC time constant also introduces time delay in addition to propagation delay of gates/buffers!

Depends on wiring resistance and total load capacitance (number of inputs connected).

CASCADED CMOS STAGES



EQUIVALENT CIRCUIT



TIMING PROBLEMS

Fast logic is time-critical!

Valid results depend on maintaining minimum overlaps (e.g. AND)
set-up times (latches)

Each logic circuit has a finite propagation delay.

This delay also depends on circuit loading, i.e. how many loads the circuit has to drive.

In addition, wiring resistance introduces delay

Depends on: number of circuits connected to wire/trace
length of trace
material

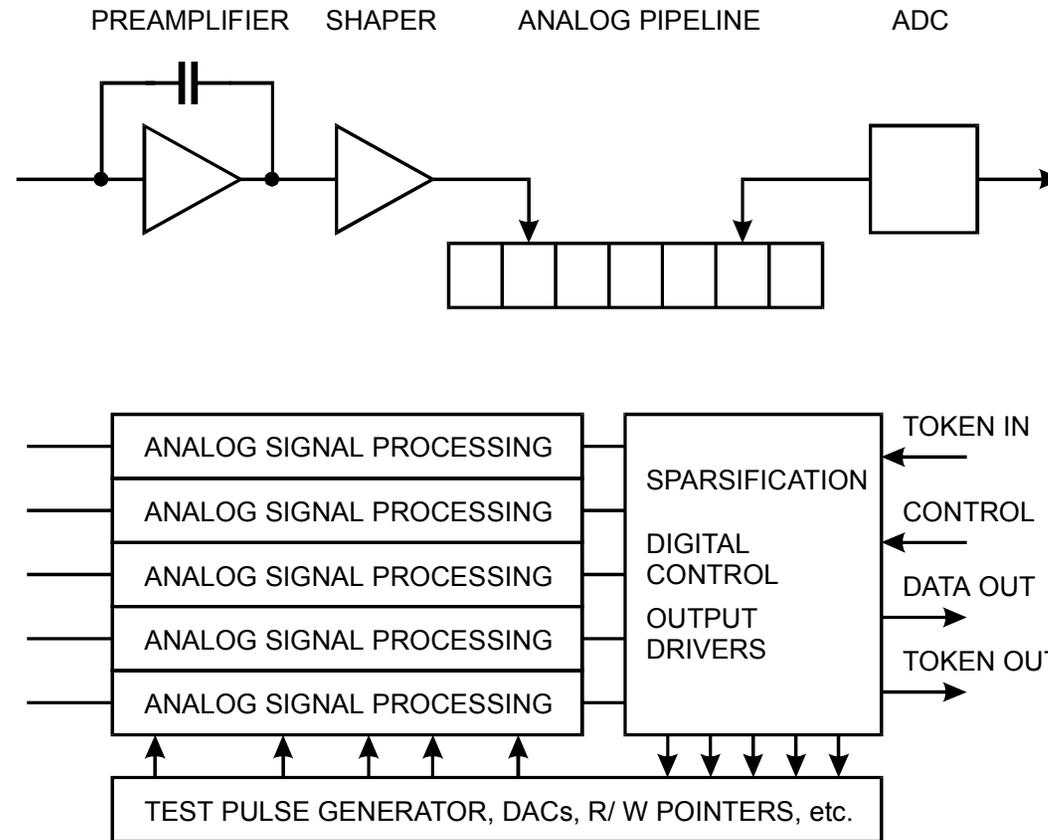
Relying on control of circuit and wiring delays to maintain timing requires great care!

Dependent on circuit variations, temperature

Most designers fail at this!

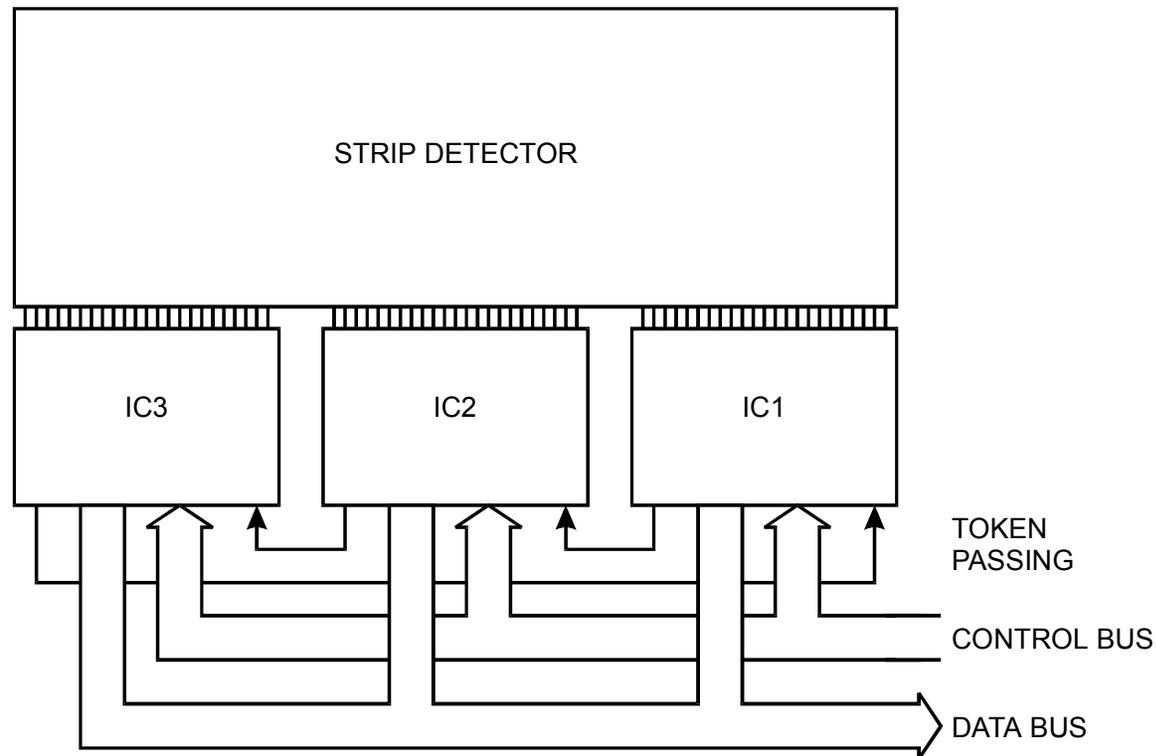
More robust: synchronous systems, where timing of all transitions is determined by master clock.

Example of readout architecture: Si strip detector



Inside a typical readout IC: 128 parallel channels of analog front-end electronics
 Logic circuitry to decode control signals, load DACs, etc.
 Digital circuitry for zero-suppression, readout

Readout of multiple ICs



IC1 is designated as master.

Readout is initiated by a trigger signal selecting appropriate time stamp to IC1.

When all data from IC1 have been transferred, a token is passed to IC2

When IC3 has finished, the token is passed back to IC1, which can begin a new cycle.

2. Digitization of Pulse Height and Time – Analog to Digital Conversion

For data storage and subsequent analysis the analog signal at the shaper output must be digitized.

Important parameters for ADCs used in detector systems:

1. Resolution
The “granularity” of the digitized output
2. Differential Non-Linearity
How uniform are the digitization increments?
3. Integral Non-Linearity
Is the digital output proportional to the analog input?
4. Conversion Time
How much time is required to convert an analog signal to a digital output?
5. Count-Rate Performance
How quickly can a new conversion commence after completion of a prior one without introducing deleterious artifacts?
6. Stability
Do the conversion parameters change with time?

Instrumentation ADCs used in industrial data acquisition and control systems share most of these requirements. However, detector systems place greater emphasis on differential non-linearity and count-rate performance. The latter is important, as detector signals often occur randomly, in contrast to measurement systems where signals are sampled at regular intervals.

1. Resolution

Digitization incurs approximation, as a continuous signal distribution is transformed into a discrete set of values. To reduce the additional errors (noise) introduced by digitization, the discrete digital steps must correspond to a sufficiently small analog increment.

Simplistic assumption:

Resolution is defined by the number of output bits, e.g. 13 bits $\rightarrow \frac{\Delta V}{V} = \frac{1}{8192} = 1.2 \cdot 10^{-4}$

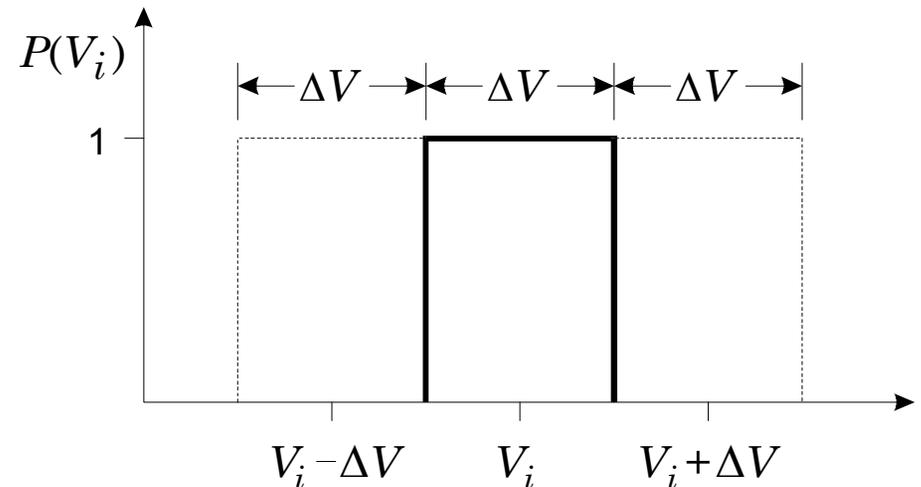
True Measure: Channel Profile

Plot probability vs. pulse amplitude that a pulse height corresponding to a specific output bin is actually converted to that address.

Ideal ADC:

Measurement accuracy:

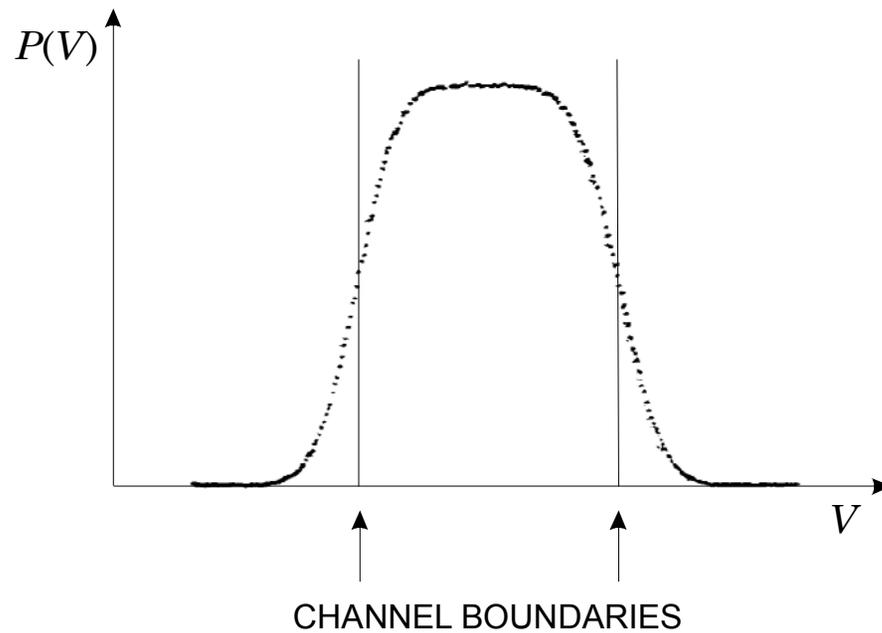
- If all counts of a peak fall in one bin, the resolution is ΔV .
- If the counts are distributed over several bins, peak fitting can yield a resolution of $10^{-1} - 10^{-2} \Delta V$,
if the distribution is known and reproducible (not necessarily a valid assumption for an ADC).



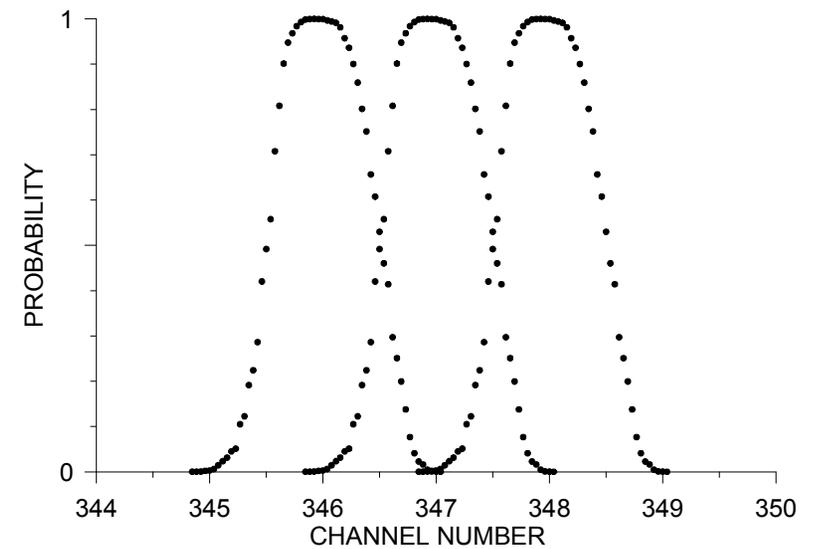
In reality, the channel profile is not rectangular as sketched above.

Electronic noise in the threshold discrimination process that determines the channel boundaries “smears” the transition from one bin to the next.

Measured channel profile (13 bit ADC)
overlap



The profiles of adjacent channels

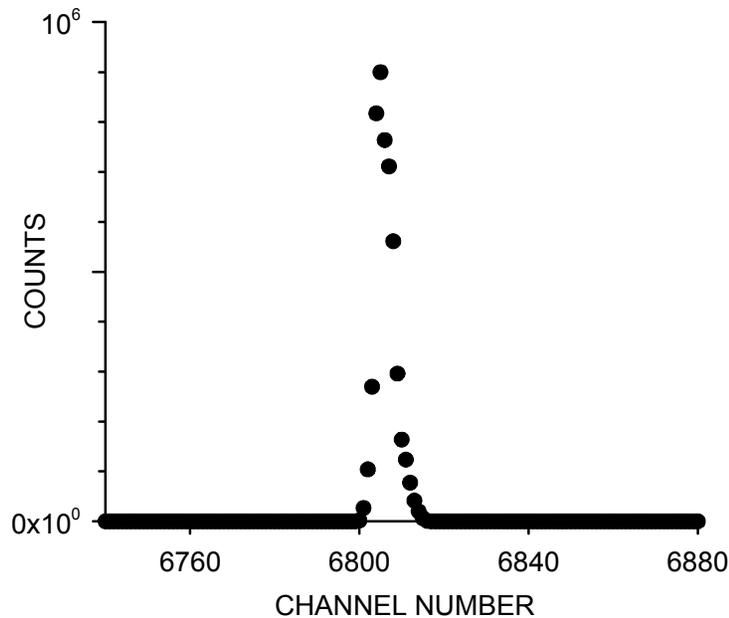


Channel profile can be checked quickly by applying the output of a precision pulser to the ADC.

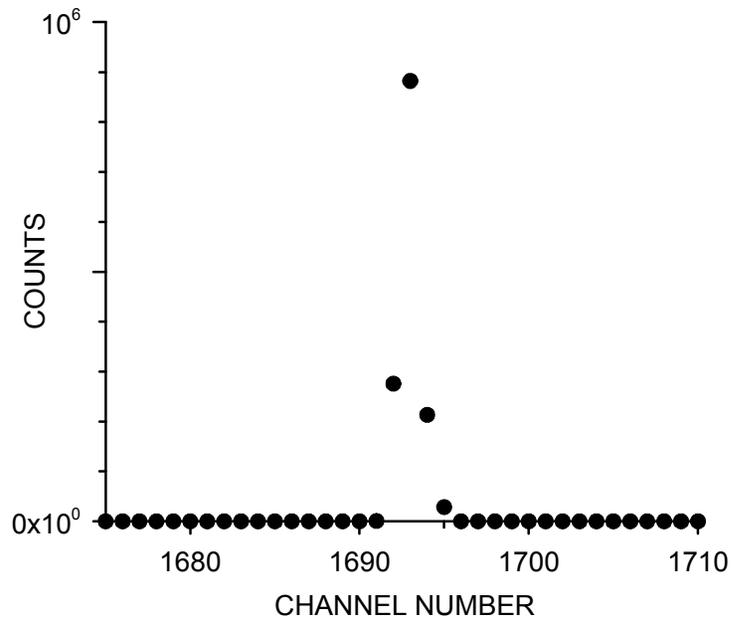
If the pulser output has very low noise, i.e. the amplitude jitter is much smaller than the voltage increment corresponding to one ADC channel or bin, all pulses will be converted to a single channel, with only a small fraction appearing in the neighbor channels.

Example of an ADC whose digital resolution is greater than its analog resolution:

8192 ch conversion range (13 bits)



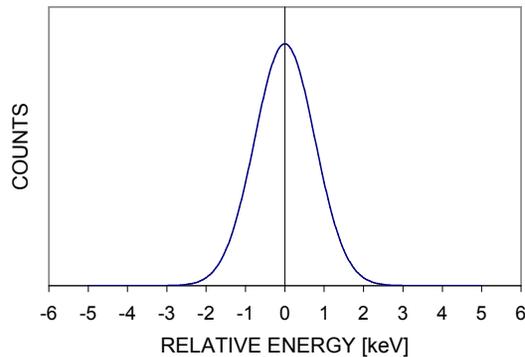
2048 ch conversion range (11 bits)



2K range provides maximum resolution – higher ranges superfluous.

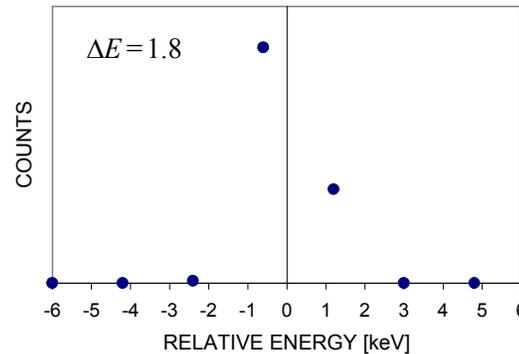
How much ADC Resolution is Required?

Example: Detector resolution ΔE
1.8 keV FWHM

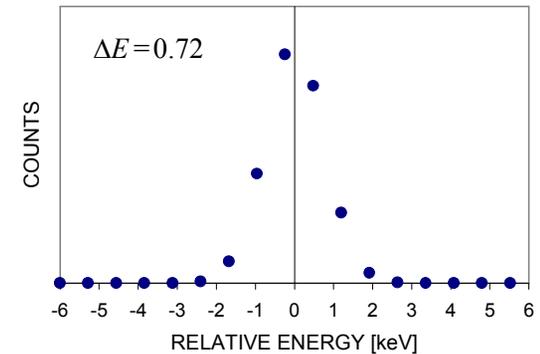


Digitized spectra for various ADC resolutions (bin widths)

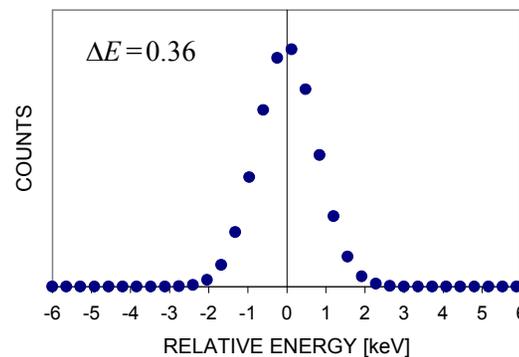
$\Delta E = 1.8 \text{ keV} = 1 \times \text{FWHM}$



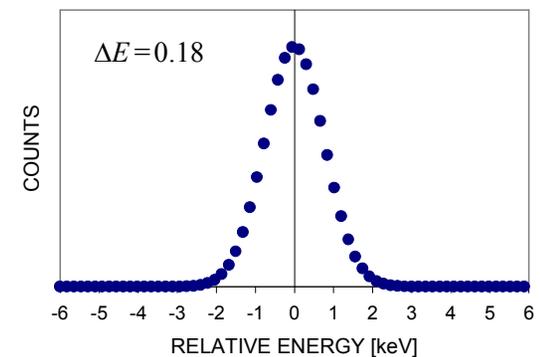
$\Delta E = 0.72 \text{ keV} = 0.4 \times \text{FWHM}$



$\Delta E = 0.36 \text{ keV} = 0.2 \times \text{FWHM}$



$\Delta E = 0.18 \text{ keV} = 0.1 \times \text{FWHM}$



Fitting can determine centroid position to fraction of bin width even with coarse digitization,
if only a single peak is present and the line shape is known.
Five digitizing channels within a linewidth (FWHM) allow robust peak finding and separation.

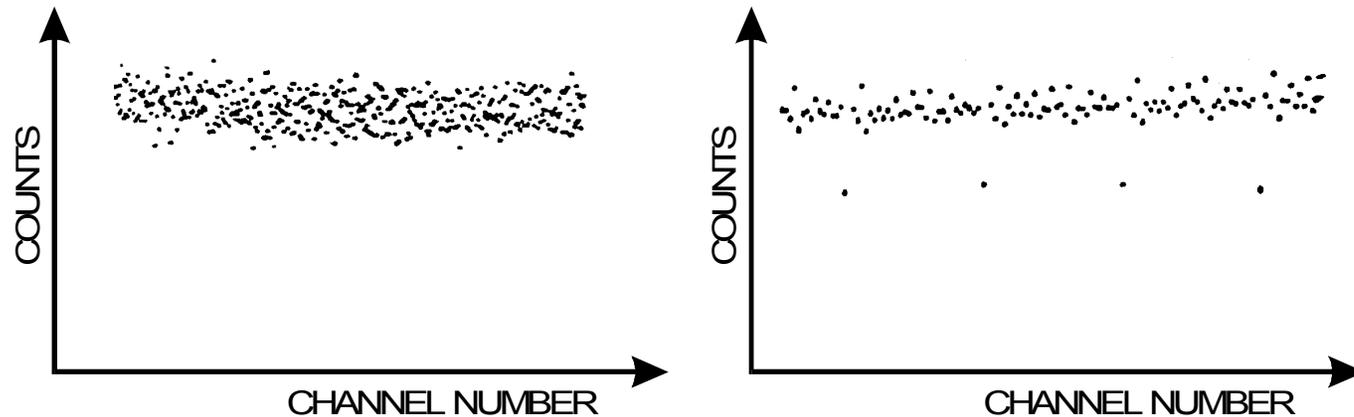
2. Differential Non-Linearity

Differential non-linearity is a measure of the inequality of channel profiles over the range of the ADC. Depending on the nature of the distribution, either a peak or an rms specification may be appropriate.

$$DNL = \max \left\{ \frac{\Delta V(i)}{\langle \Delta V \rangle} - 1 \right\} \quad \text{or} \quad DNL = \text{r.m.s.} \left\{ \frac{\Delta V(i)}{\langle \Delta V \rangle} - 1 \right\}$$

where $\langle \Delta V \rangle$ is the average channel width and $\Delta V(i)$ is the width of an individual channel.

Typical differential non-linearity patterns (“white” input spectrum). vertical scale 10x larger than left
↓



An ideal ADC would show an equal number of counts in each bin.
The spectrum to the left shows a random pattern,
but note the multiple periodicities visible in the right hand spectrum.

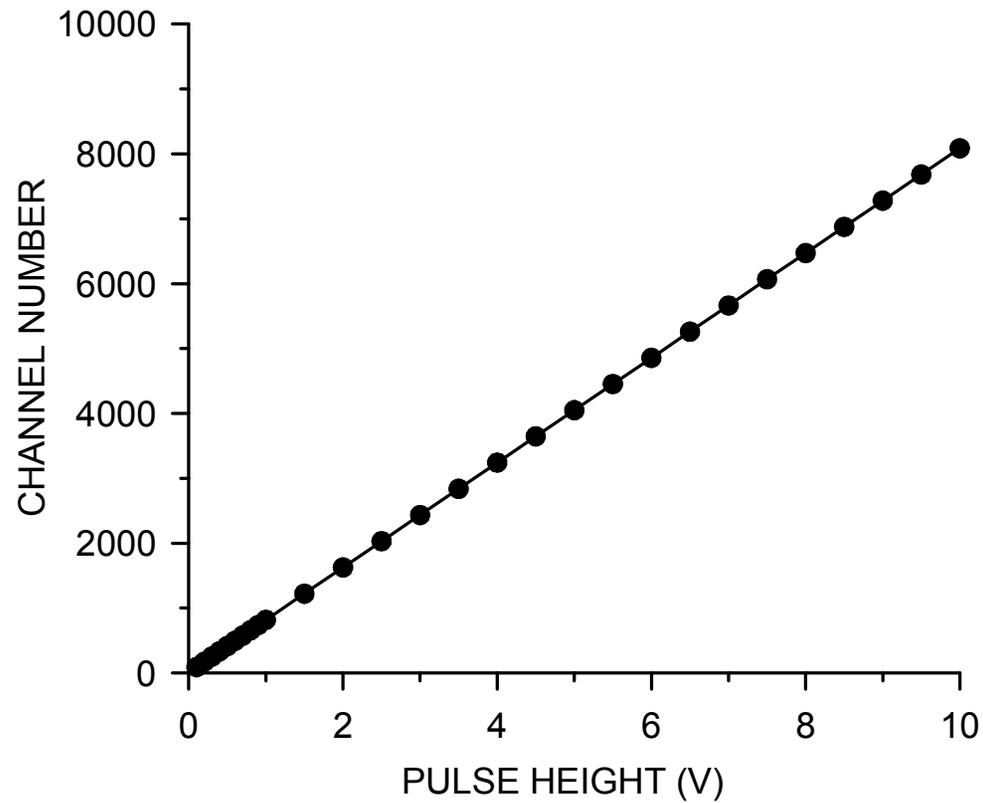
Differential non-linearity of $< \pm 1\%$ max. is typical, but state-of-the-art ADCs can achieve 10^{-3} rms, i.e. the variation is comparable to the statistical fluctuation for 10^6 random counts.

Note: Instrumentation ADCs are often specified with an accuracy of ± 0.5 LSB (least significant bit) or more, so

1. the differential non-linearity may be 50% or more,
 2. the response may be non-monotonic
- ⇒ output may decrease when input rises.

3. Integral Non-Linearity

Integral non-linearity measures the deviation from proportionality of the measured amplitude to the input signal level.



The dots are measured values and the line is a fit to the data.

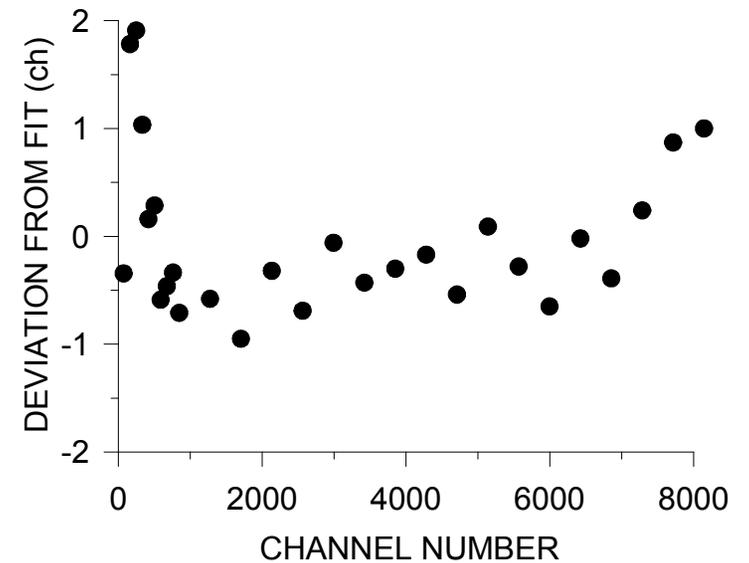
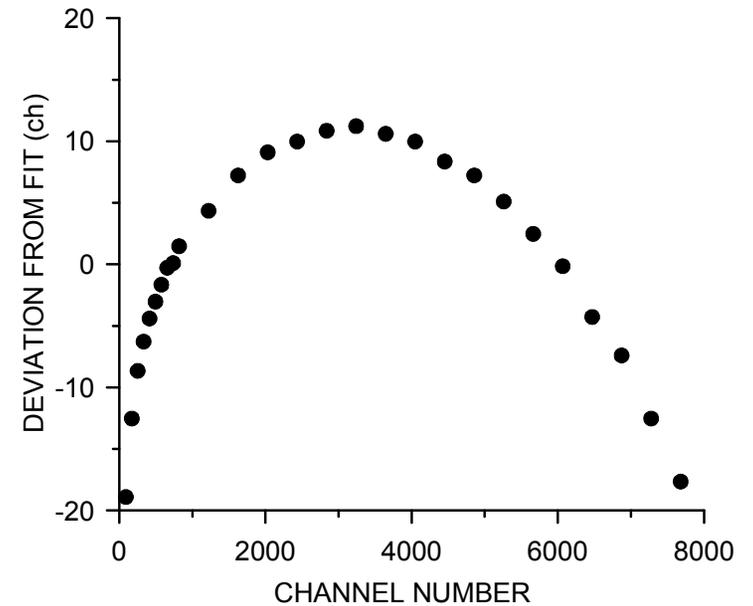
This plot is not very useful if the deviations from linearity are small.

Plotting the deviations of the measured points from the fit yields a more useful result.

Integral non-linearity measured with a 400 ns wide input pulse

The linearity of an ADC can depend on the input pulse shape and duration, due to bandwidth limitations in the circuitry.

Increasing the pulse width to 3 μ s improved the linearity significantly:



4. Conversion Time

During the acquisition of a signal the system cannot accept a subsequent signal (“dead time”)

Dead Time =

- signal acquisition time → time-to-peak + const.
- + conversion time → can depend on pulse height
- + readout time to memory → depends on speed of data transmission and buffer memory access

Dead time affects measurements of yields or reaction cross-sections. Unless the event rate $\ll 1/(\text{dead time})$, it is necessary to measure the dead time, e.g. with a reference pulser fed simultaneously into the spectrum.

The total number of reference pulses issued during the measurement is determined by a scaler and compared with the number of pulses recorded in the spectrum.

Does a pulse-height dependent dead time mean that the correction is a function of pulse height?

Usually not. If events in different part of the spectrum are not correlated in time, i.e. random, they are all subject to the same average dead time (although this average will depend on the spectral distribution).

- Caution with correlated events!
 Example: Decay chains, where lifetime is $<$ dead time.
 The daughter decay will be lost systematically.

5. Count Rate Effects

Problems are usually due to internal baseline shifts with event rate or undershoots following a pulse.

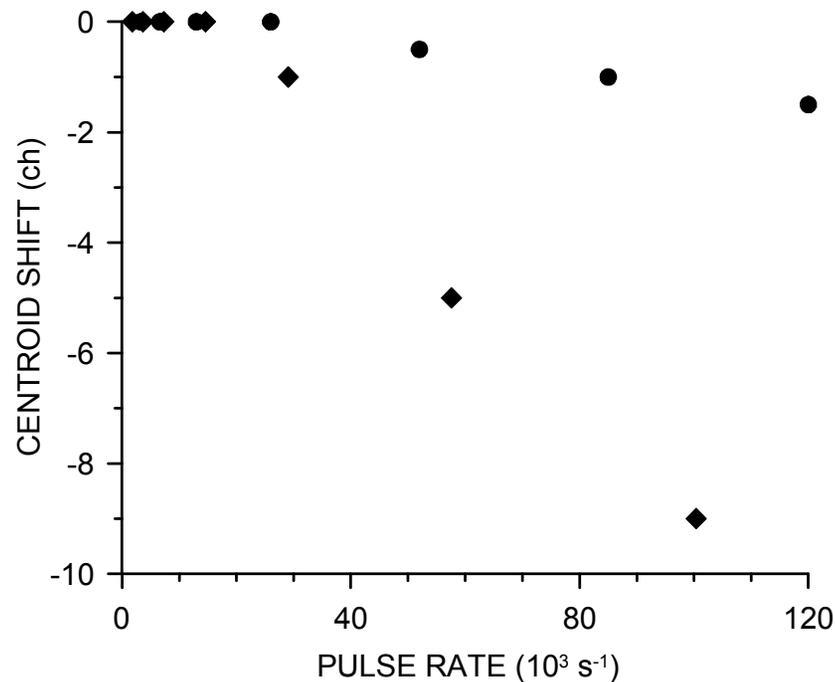
If signals occur at constant intervals, the effect of an undershoot will always be the same.

However, in a random sequence of pulses, the effect will vary from pulse to pulse.

⇒ spectral broadening

Baseline shifts tend to manifest themselves as a systematic shift in centroid position with event rate.

Centroid shifts for two 13 bit ADCs vs. random rate:



6. Stability

Stability vs. temperature is usually adequate with modern electronics in a laboratory environment.

- Note that temperature changes within a module are typically much smaller than ambient.

However: Highly precise or long-term measurements require spectrum stabilization to compensate for changes in gain and baseline of the overall system.

Technique: Using precision pulsers place a reference peak at both the low and high end of the spectrum.

(Pk. Pos. 2) – (Pk. Pos. 1) → Gain, ...

then

(Pk. Pos. 1) or (Pk. Pos. 2) → Offset

Traditional Implementation: Hardware,
spectrum stabilizer module

Today, it is more convenient to determine the corrections in software.

These can be applied to calibration corrections or used to derive an electrical signal that is applied to the hardware (simplest and best in the ADC).

Analog to Digital Conversion Techniques

1. Flash ADC

The input signal is applied to n comparators in parallel. The switching thresholds are set by a resistor chain, such that the voltage difference between individual taps is equal to the desired measurement resolution.

2^n comparators for n bits (8 bit resolution requires 256 comparators)

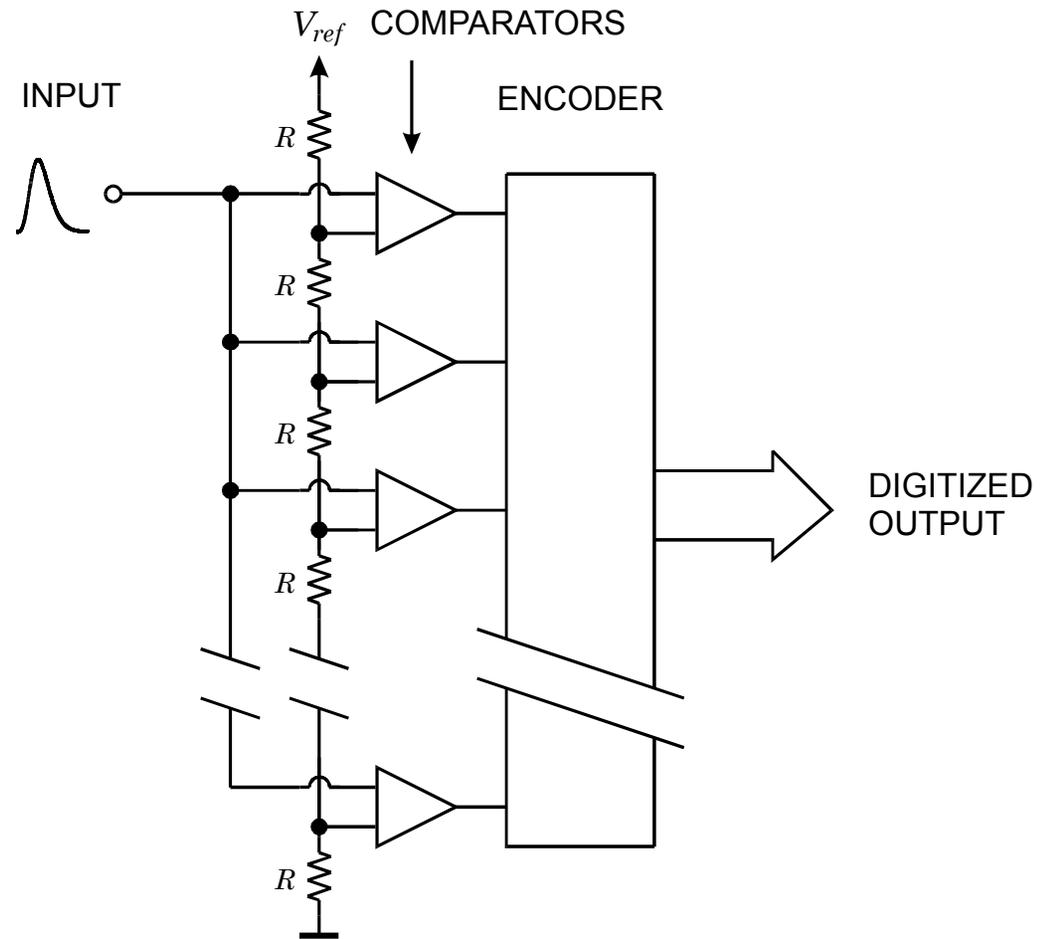
Feasible in monolithic ICs since the absolute value of the resistors in the reference divider chain is not critical, only the relative matching.

Advantage: short conversion time
(<10 ns available)

Drawbacks:

- limited accuracy
- (many comparators)
- power consumption
- Differential non-linearity $\sim 1\%$
- High input capacitance

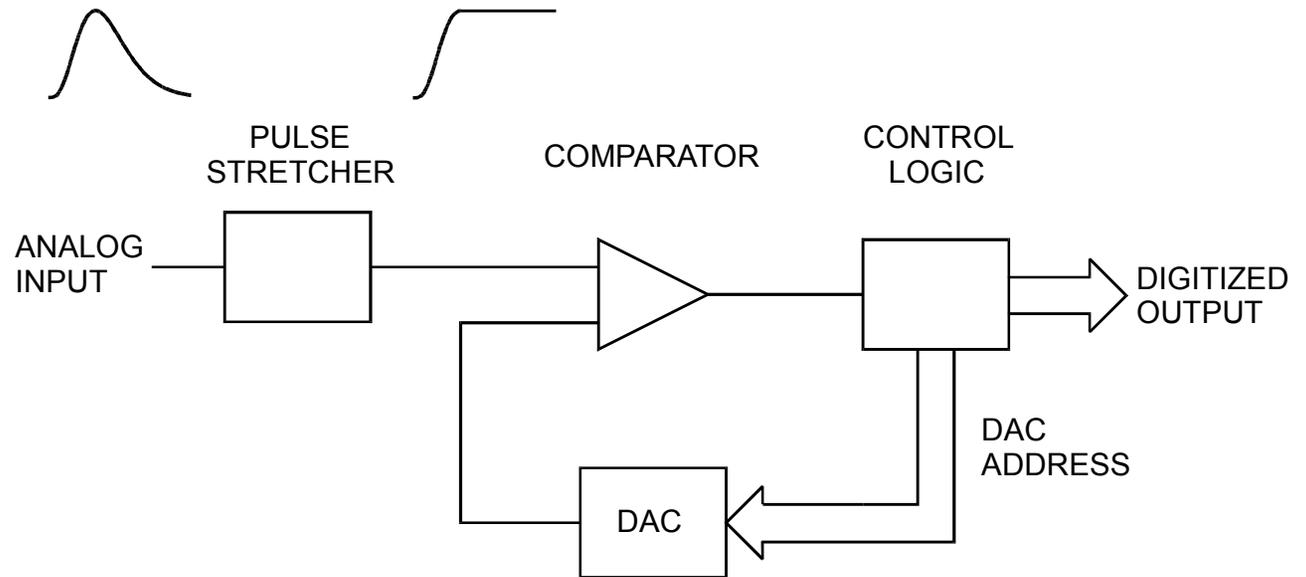
Speed is often limited by the analog driver feeding the input.



2. Successive Approximation ADC

Sequentially raise comparator threshold proportional to $2^n, 2^{n-1}, \dots, 2^0$ and set corresponding bit if the comparator output is high (DAC output < pulse height)

n conversion steps yield 2^n channels, i.e. 8K channels require 13 steps



Advantages: speed ($\sim \mu\text{s}$)
 high resolution
 ICs (monolithic + hybrid) available

Drawback: Differential non-linearity (typ. 10 – 20%)

Reason: Resistors that set DAC output must be extremely accurate. For $\text{DNL} < 1\%$ the resistor determining the 2^{12} level in an 8K ADC must be accurate to $< 2.4 \cdot 10^{-6}$.

DNL can be corrected by various techniques:

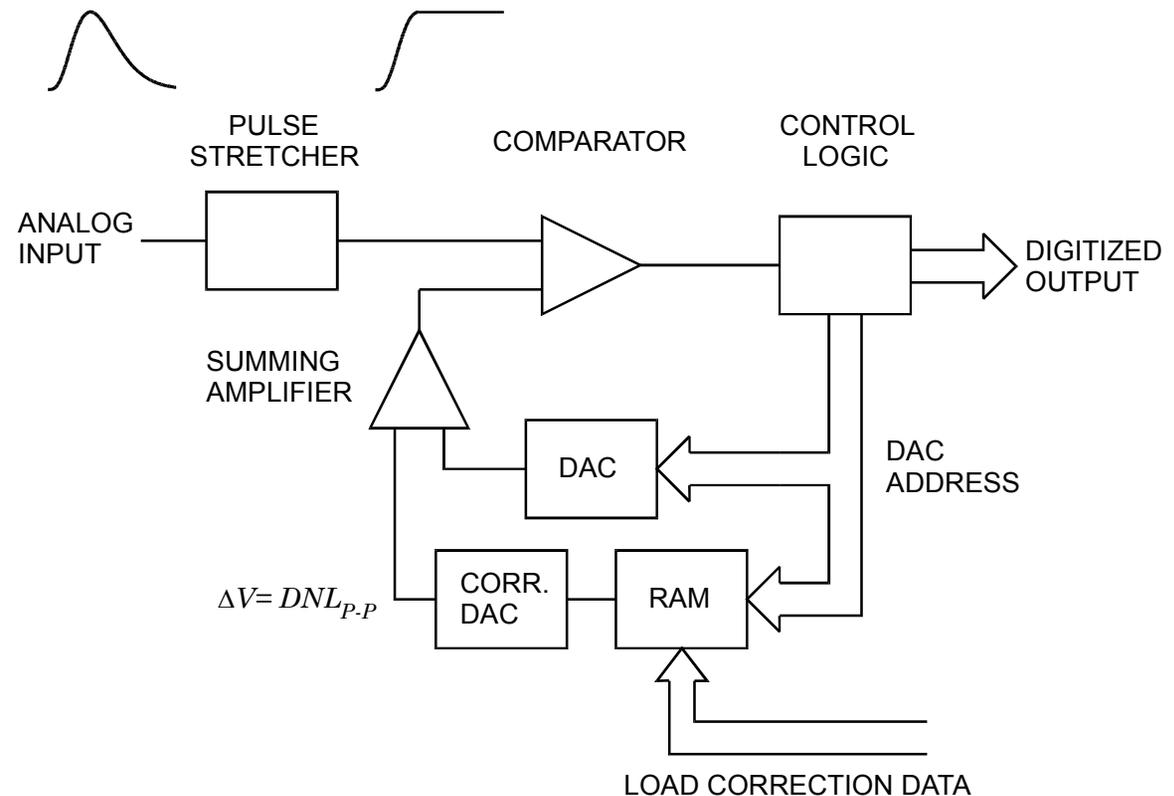
- averaging over many channel profiles for a given pulse amplitude (“sliding scale” or “Gatti principle”)
- correction DAC (“brute force” application of IC technology)

The primary DAC output is adjusted by the output of a correction DAC to reduce differential non-linearity.

Correction data are derived from a measurement of DNL. Corrections for each bit are loaded into the RAM, which acts as a look-up table to provide the appropriate value to the correction DAC for each bit of the main DAC.

The range of the correction DAC must exceed the peak-to-peak differential non-linearity.

If the correction DAC has N bits, the maximum DNL is reduced by $1/2^{(N-1)}$ (if deviations are symmetrical).

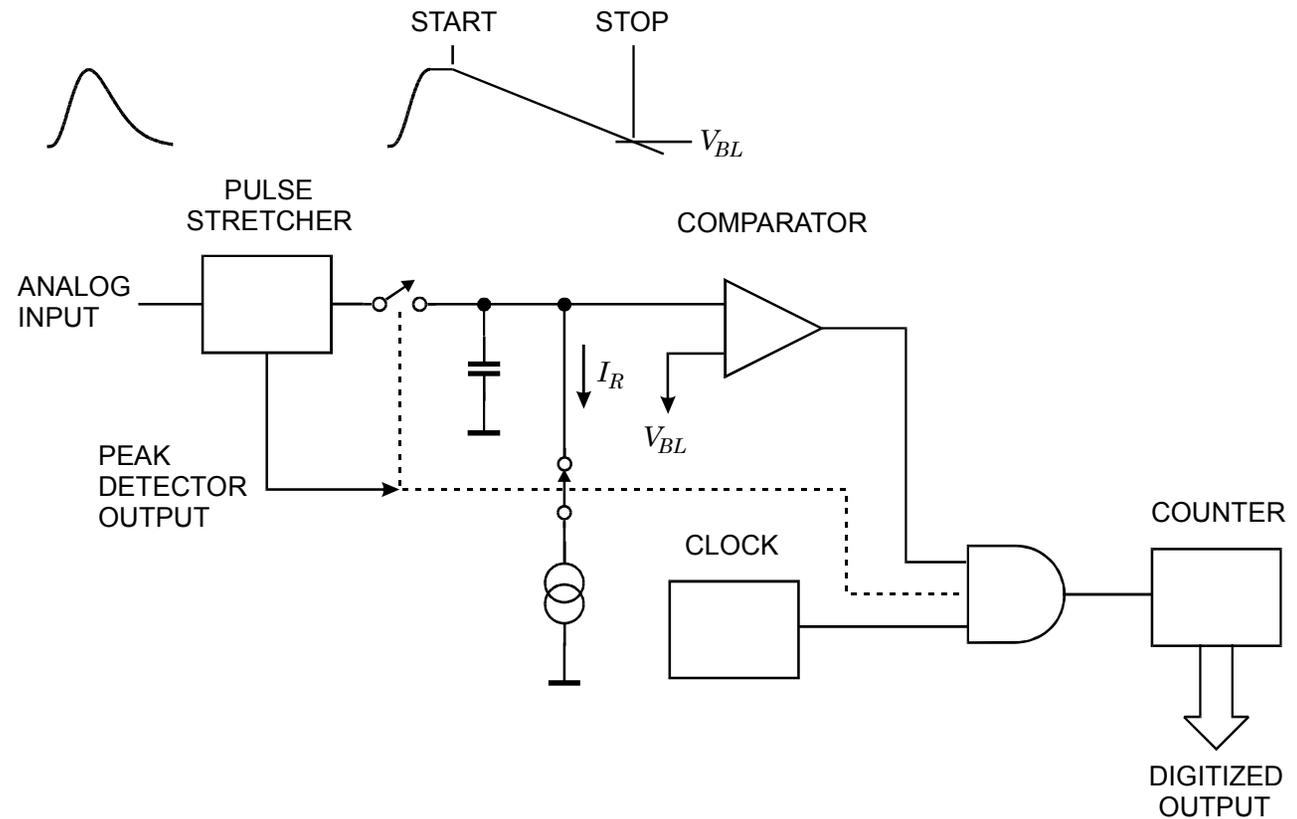


3. Wilkinson ADC

The peak signal amplitude is acquired by a pulse stretcher and transferred to a memory capacitor.

Then, simultaneously,

1. the capacitor is disconnected from the stretcher,
2. a current source is switched on to linearly discharge the capacitor,
3. a counter is enabled to determine the number of clock pulses until the voltage on the capacitor reaches the baseline.



Advantage: excellent differential linearity (continuous conversion process)

Drawbacks: slow – conversion time = $n \cdot T_{clock}$ (n = channel number \propto pulse height)

$$T_{clock} = 10 \text{ ns} \rightarrow T_{conv} = 82 \mu\text{s} \text{ for 13 bits}$$

Clock frequencies of 100 MHz typical,
but >400 MHz possible with excellent performance

“Standard” technique for high-resolution spectroscopy.

4. Pipelined ADCs

Most common architecture for high-speed high-resolution ADCs

Input to each stage is fed both to a sample and hold (S&H) and a 3-bit flash ADC.

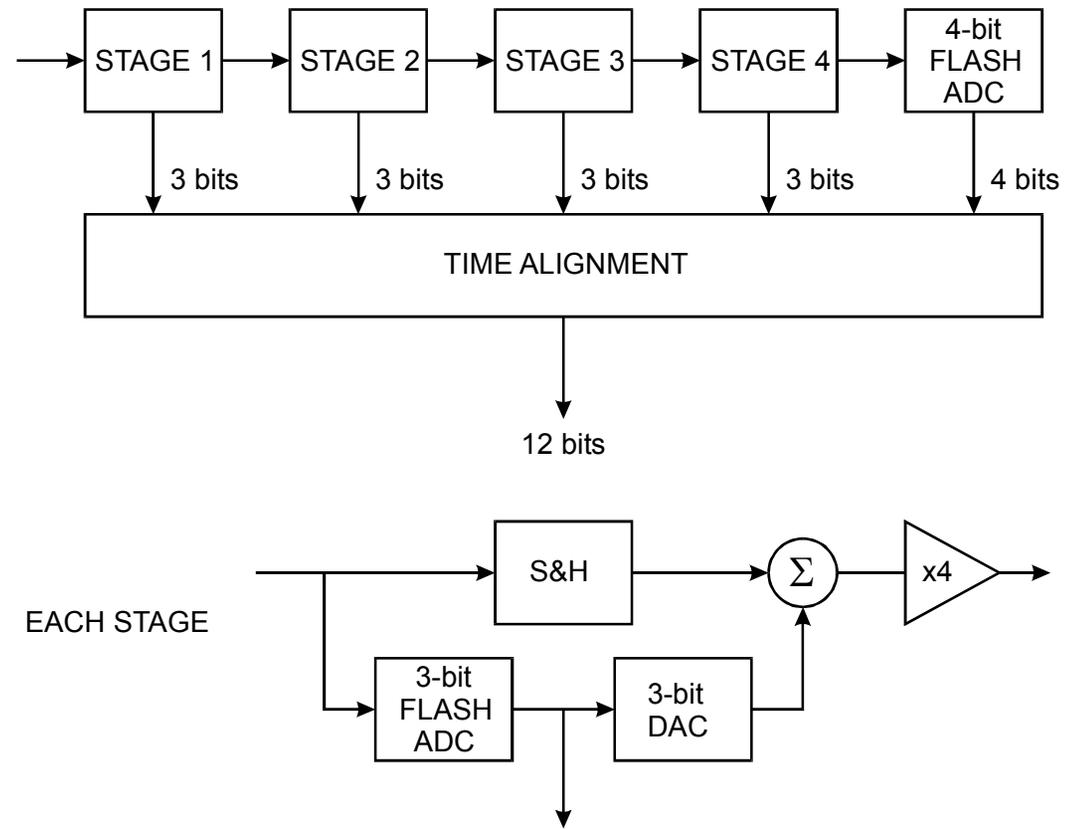
The S&H maintains the signal level during conversion. The flash ADC quantizes its input to 3 bit accuracy. This output is fed to a DAC with 12 bit accuracy. The DAC's analog output is subtracted from the original signal and the difference signal is passed on to the next stage.

The last 4 bits are resolved by a 4-bit flash ADC.

As soon as a stage has passed its result to the next stage it can begin processing the next signal, so throughput is not determined by the total conversion time, but by the time per stage.

Since the interstage gain is only 4 (rather than 8 corresponding to 3 bits), each stage only contributes 2 bits of resolution. The extra bit is used for error correction.

Commercially available: 1 GS/s conversion rates with 8-bit resolution and a power dissipation of about 1.5 W.



Hybrid Analog-to-Digital Converters

Conversion techniques can be combined to obtain high resolution and short conversion time.

1. Flash + Successive Approximation or Flash + Wilkinson (Ramp Run-Down)

Utilize fast flash ADC for coarse conversion (e.g. 8 out of 13 bits)

Successive approximation or Wilkinson converter to provide fine resolution.

Limited range, so short conversion time: 256 ch with 100 MHz clock \Rightarrow 2.6 μ s

Results: 13 bit conversion in $< 4 \mu$ s with excellent integral and differential linearity

2. Flash ADCs with Sub-Ranging

Not all applications require constant absolute resolution over the full range. Sometimes only *relative* resolution must be maintained, especially in systems with a very large dynamic range.

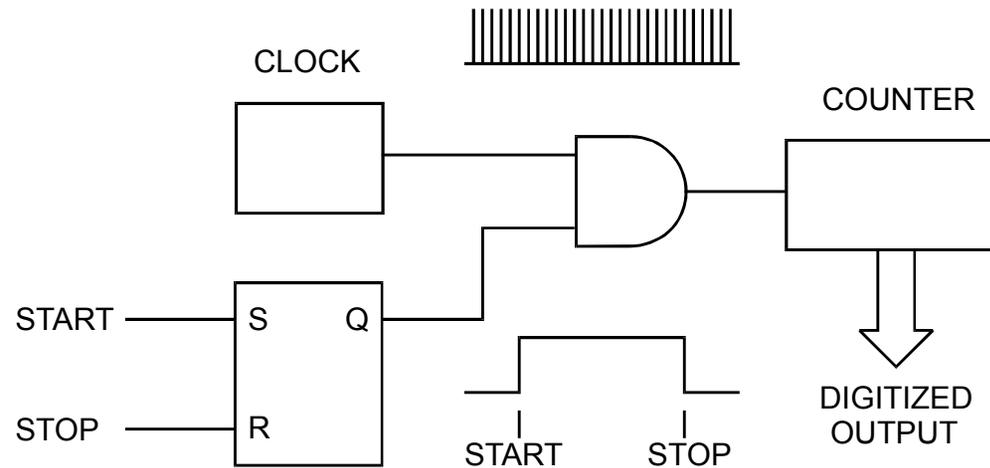
Precision binary divider at input to determine coarse range + fast flash ADC for fine digitization.

Example: Fast digitizer that fits in phototube base (FNAL)

17 to 18 bit dynamic range
Digital floating point output
(4 bit exponent, 8+1 bit mantissa)
16 ns conversion time

Time Digitizers

1. Counter



Simplest arrangement: Count clock pulses between start and stop.

Limitation: Speed of counter

Current technology limits speed of counter system to about 1 GHz

$$\Rightarrow \Delta t = 1 \text{ ns}$$

Advantages: Simplicity
Multi-hit capability

2. Analog Ramp

Commonly used in high-resolution digitizers ($\Delta t = 10$ ps)

Principle:

Charge capacitor through switchable current source

Start pulse: turn on current source

Stop pulse: turn off current source

⇒ Voltage on storage capacitor

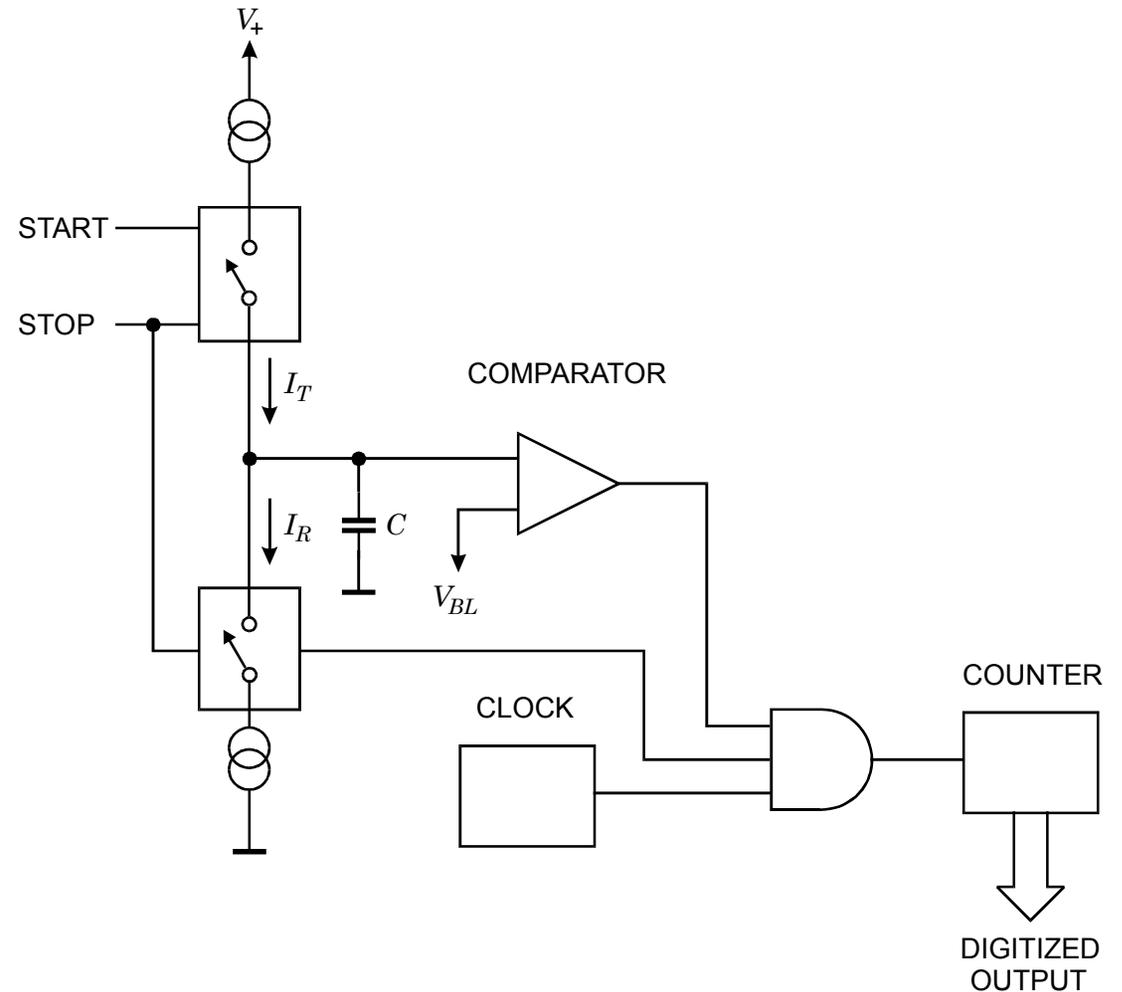
Use Wilkinson ADC with smaller discharge current to digitize voltage.

Drawbacks: No multi-hit capability

Deadtime

Advantages: High resolution (~ps)

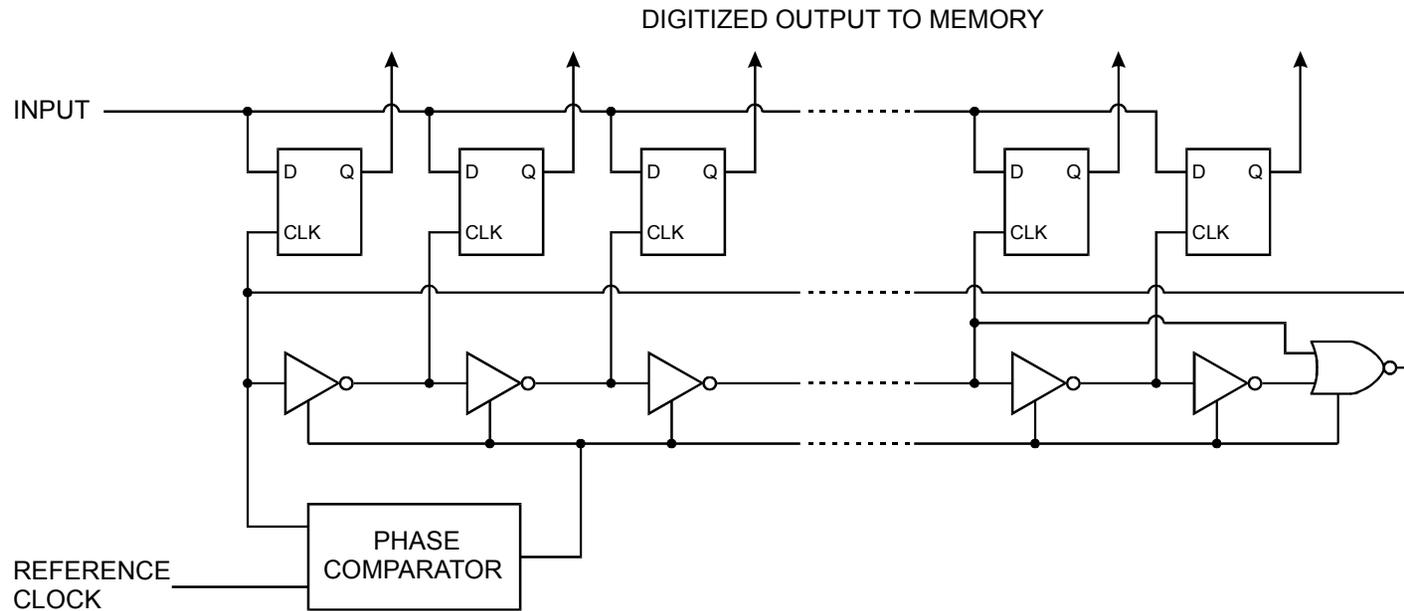
Excellent differential linearity



3. Digitizers with Clock Interpolation

Most experiments in HEP require multi-hit capability, no deadtime

Commonly used in HEP ICs for time digitization (Y. Arai, KEK)



Clock period interpolated by inverter delays (U1, U2, ...).

Delay can be fine-tuned by adjusting operating point of inverters. Stabilized by delay locked loop.

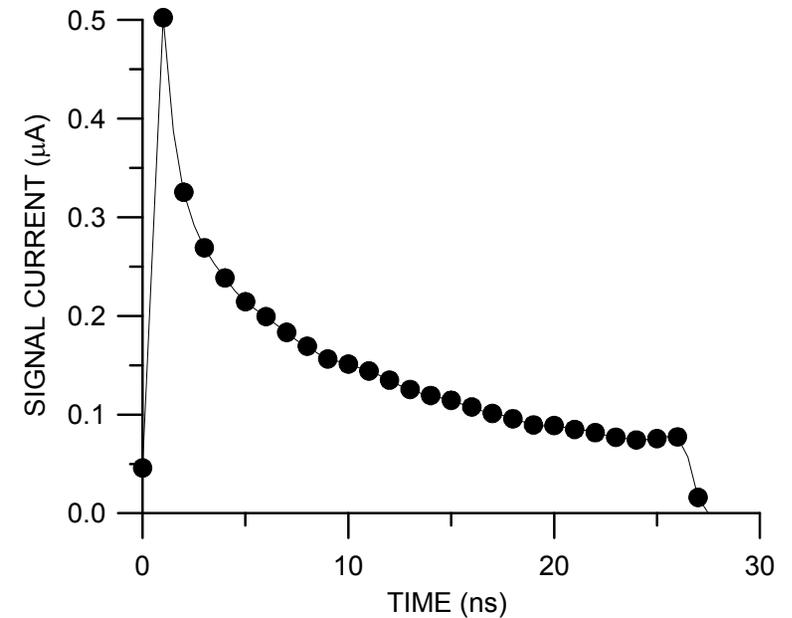
Devices with 250 ps resolution fabricated and tested.

see Y. Arai et al., IEEE Trans. Nucl. Sci. **NS-45/3** (1998) 735-739 and references therein.

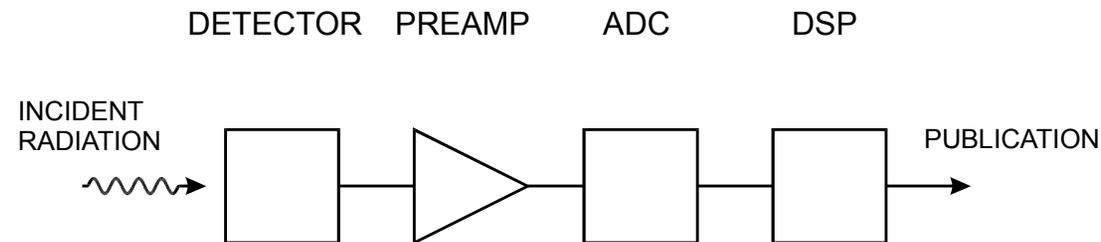
3. Digital Signal Processing

Sample detector signal with fast digitizer to reconstruct pulse:

Then use digital signal processor to perform mathematical operations for desired pulse shaping.



Block Diagram

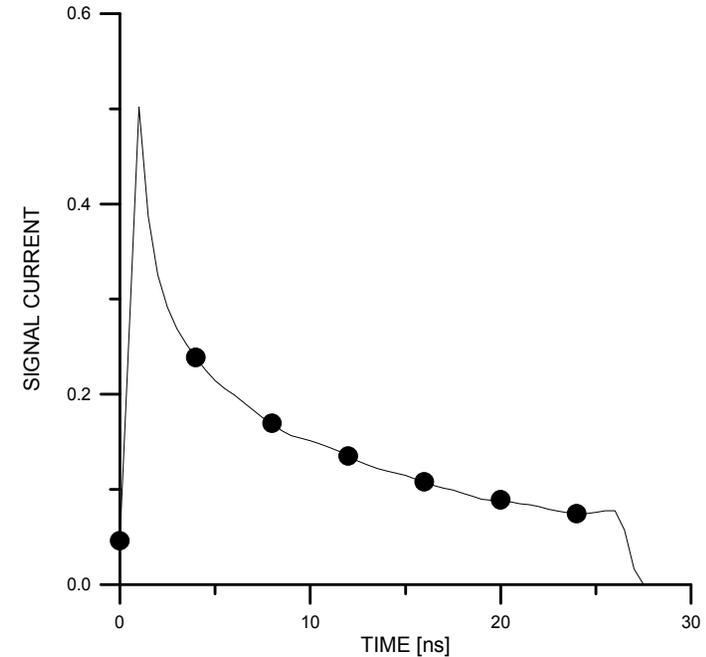


DSP allows great flexibility in implementing filtering functions

However: increased circuit complexity
increased demands on ADC, compared to traditional shaping.

Important to choose sample interval sufficiently small to capture pulse structure.

Sampling interval of 4 ns misses initial peak.

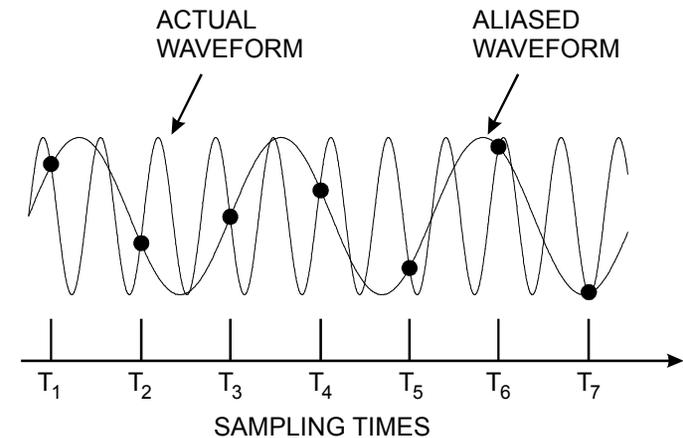


With too low a sampling rate high frequency components will be “aliased” to lower frequencies:

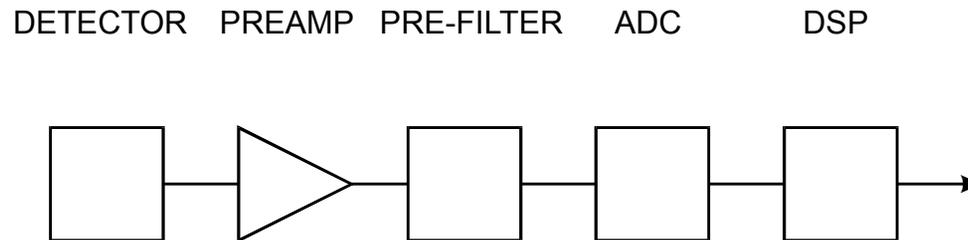
Applies to any form of sampling
(time waveform, image, ...)

Nyquist condition:

Sampling frequency > 2x highest signal frequency



⇒ Fast ADC required + Pre-Filter to limit signal bandwidth



- Dynamic range requirements for ADC may be more severe than in analog filtered system (depending on pulse shape and pre-filter).
- Digitization introduces additional noise (“quantization noise”)

If one bit corresponds to an amplitude interval Δ , the quantization noise

$$\sigma_v^2 = \int_{-\Delta/2}^{\Delta/2} \frac{v^2}{\Delta} dv = \frac{\Delta^2}{12} .$$

(differential non-linearity introduces quasi-random noise)

- Electronics preceding ADC and front-end of ADC must exhibit same precision as analog system, i.e. baseline and other pulse-to-pulse amplitude fluctuations less than order $Q_n/10$, i.e. typically 10^{-4} in high-resolution systems. For 10 V FS at the ADC input this corresponds to < 1 mV.

⇒ ADC must provide high performance at short conversion times. Today this is technically feasible for some applications, e.g. detectors with moderate to long collection times (γ and x-ray detectors).

Digital Filtering

Filtering is performed by convolution:

$$S_o(n) = \sum_{k=0}^{N-1} W(k) \cdot S_i(n - k)$$

$W(k)$ is a set of coefficients that describes the weighting function yielding the desired pulse shape.

A filter performing this function is called a Finite Impulse Response (FIR) filter.

This is analogous to filtering in the frequency domain.

In the frequency domain the result of filtering is determined by multiplying the responses of the individual stages:

$$G(f) = G_1(f) \cdot G_2(f)$$

where $G_1(f)$ and $G_2(f)$ are complex numbers.

The theory of Fourier transforms states that the equivalent result in the time domain is formed by convolution of the individual time responses:

$$g(t) = g_1(t) * g_2(t) \equiv \int_{-\infty}^{+\infty} g_1(\tau) \cdot g_2(t - \tau) d\tau ,$$

analogously to the discrete sum shown above.

Benefits of digital signal processing:

- Flexibility in implementing filter functions
- Filters possible that are impractical in hardware
- Simple to change filter parameters
- Tail cancellation and pile-up rejection easily incorporated
- Adaptive filtering can be used to compensate for pulse shape variations.

Where is digital signal processing appropriate?

- Systems highly optimized for
Resolution
High counting rates
- Variable detector pulse shapes

Where is analog signal processing best (most efficient)?

- Fast shaping
- Systems not sensitive to pulse shape (fixed shaper constants)
- High density systems that require small circuit area or low power

Both types of systems require careful analog design.

Progress in fast ADCs (precision, reduced power) will expand range of DSP applications